Karl-Franzens-University Graz

Faculty of Social and Economic Sciences

A High-Frequency Analysis of Bitcoin Markets

Alexander Brauneis, Roland Mestel, Ryan Riordan, Erik Theissen

Working Paper 2018-06

October 25, 2018

Working Paper Series
Faculty of Social and Economic Sciences
Karl-Franzens-University Graz
ISSN 2304-7658
sowi.uni-graz.at/forschung/working-paper-series/
sowi-wp@uni-graz.at

UNIVERSITY OF GRAZ

A High-Frequency Analysis of Bitcoin Markets

Alexander Brauneis[a], Roland Mestel[b,], Ryan Riordan[c], Erik Theissen[d,b]*

Working Paper 2018-06

October 25, 2018

Abstract

We study Bitcoin (BTC) exchanges on three continents, Bitfinex, Bitstamp and GDAX. We use a high frequency dataset that contains both transactions data and snapshots of the BTC to US dollar (BTCUSD) order book. The BTCUSD market is highly liquid in terms of bid-ask spreads and order book depth. While spreads are low, we find large differences between the three exchanges in terms of transaction and posted prices. The price differences fall over our sample period meaning that markets are becoming more integrated. We show that exchanges play an increasingly important role in the transfer of BTC. At the end of 2017, exchanges processed roughly 30% of BTC transfers at the end of our sample period this increases to 90%.

Keywords: bitcoin, cryptocurrencies, liquidity, bid-ask spread

JEL: G10, G12

[a] University of Klagenfurt, Department of Finance, Universitaetsstrasse 65-67, 9020 Klagenfurt, AUSTRIA

[b] University of Graz, Department of Banking and Finance, Universitaetsstrasse 15, 8010 Graz, AUSTRIA

[c] Queens University, Smith School of Business, 143 Union Street West, Kingston ON K7L 3N6, CANADA

[d] University of Mannheim, Department of Finance, L9, 1-2, 68161 Mannheim, GERMANY

[*] Corresponding author. Tel.: +43(316)380-7304, E-Mail roland.mestel @uni-graz.at.

Working Paper Series
Faculty of Social and Economic Sciences
Karl-Franzens-University Graz
ISSN 2304-7658
sowi.uni-graz.at/forschung/working-paper-series/
sowi-wp@uni-graz.at

UNIVERSITY OF GRAZ

A High-Frequency Analysis of Bitcoin Markets

Alexander Brauneis*, Roland Mestel†, Ryan Riordan‡, Erik Theissen§

This version: October 24, 2018

Abstract: We study Bitcoin (BTC) exchanges on three continents, Bitfinex, Bitstamp and GDAX. We use a high frequency dataset that contains both transactions data and snapshots of the BTC to US dollar (BTCUSD) order book. The BTCUSD market is highly liquid in terms of bid-ask spreads and order book depth. While spreads are low, we find large differences between the three exchanges in terms of transaction and posted prices. The price differences fall over our sample period meaning that markets are becoming more integrated. We show that exchanges play an increasingly important role in the transfer of BTC. At the end of 2017, exchanges processed roughly 30% of BTC transfers at the end of our sample period this increases to 90%.

Keywords: Bitcoin, Cryptocurrencies, Liquidity, Bid-Ask Spread
JEL Classifications: G10, G12

*University of Klagenfurt, e-mail: *alexander.brauneis@aau.at.*
†University of Graz, e-mail: *roland.mestel@uni-graz.at.*
‡Queens University, e-mail: *ryan.riordan@queensu.ca.*
§University of Mannheim and University of Graz, e-mail: *theissen@uni-mannheim.de.*

1 Introduction

Cryptocurrencies are a new class of digital assets that can be held in virtual wallets, on exchanges, and in retirement accounts. Bitcoin is the most prominent cryptocurrency and was introduced by Satoshi Nakamoto in 2008 (Nakamoto 2008). Bitcoin is a digital asset issued periodically (on average every 10 minutes) that can be exchanged for other currencies, fiat or crypto, at one of many crypto-exchanges. The market capitalization of Bitcoin (BTC) peaked at USD 336 billions in Dec. 2017 with an all-time high of 20,089 BTCUSD (see coinmarketcap.com). Most research has focused on the Blockchain (Abadi and Brunnermeier 2018, Cong and He 2018), the ledger behind BTC, or on explaining the meteoric price increase of Bitcoin (Gandal et al. 2018). Less research has focused on the crypto-exchanges that determine the exchange rate between BTC and US Dollars, Euros, or Japanese Yen. These exchanges are important as they allow participants to exchange Bitcoin for currencies that can be more ubiquitously used to purchase goods and services, and to pay taxes. The exchanges also produce reference prices for the settlement of futures, options, and exchange-traded-funds. Finally other crytocurrencies are "priced" off of BTC meaning that how the price of BTC is determined and how the exchanges operate is an important component of all cryptocurrencies.

This paper contributes to the emerging academic literature on cryptocurrencies by providing a high-frequency analysis of the BTCUSD market. Bitcoin exchanges, like most equity exchanges, operate centralized limit order books and trading is fragmented across a number of trading venues. Otherwise trading Bitcoin is different from trading other financial assets in many potentially important ways. Trading occurs around the globe 24 hours a day, 365 days a year. All traders have direct access to the exchange instead of accessing it through a broker. Exchanges allow direct transfer from and to bank accounts or credit cards. Transactions are cleared and settled by exchanges directly. Traders have a fiat currency and a Bitcoin account that are administered by the exchange. Trades are settled by transferring fiat currency and Bitcoins between these accounts. Exchanges allow margin trading and short-selling in very limited circumstances. Regulatory oversight is low and uncertain. Understanding the market structure and how prices evolve are important to understand how unregulated and decentralized markets develop.

We study market integration, liquidity, and the market structure of BTCUSD trading on three leading cryptocurrency exchanges, Bitfinex, Bitstamp, and GDAX for the heavily traded currency pair BTCUSD. Bitfinex, Bitstamp and GDAX, are located in Asia, Europe and the US respectively and trade 650 million USD daily and represent roughly 74% of the total BTCUSD trading volume in our sample period (source: data.bitcoinity.org, accessed: June 15, 2018). Consistent with Kroeger and Sarkar 2017, our results show that BTC prices across exchanges are not well integrated. In a price series of BTC transactions from 2015 to 2018, large and long-lasting price differences persist between the exchanges. This suggests that arbitrage opportunities are not always exploited and prices are inefficient in that they do not accurately represent aggregate demand and supply. The large price discrepancies may exist for a number of reasons. For instance, the cost to transfer fiat or BTC from one exchange to another exchange may be too costly to make exploiting the arbitrage opportunity worthwhile. Some market participants may be restricted to trading on a market in their jurisdiction. Transferring BTC between markets can take at least 10 minutes and considerably longer depending on Blockchain congestion. Setting up multiple accounts is time-consuming. The economics literature refers to these problems as frictions and they can hinder market functioning (Shleifer and Vishny 1997).

The large violations of the law of one price make a statistical analysis of price integration difficult. In contrast to Kroeger and Sarkar 2017, we use high frequency order book data to study the BTCUSD market. We obtain over 5 mio order book snapshots from December 15th 2017 to June 15th 2018. The order book data allow for a granular analysis that does not rely solely on transaction data. Most financial markets rely on the bid price (the price at which traders can instantly sell) being lower than the ask price (the price at which traders can instantly buy). When markets are fragmented and trading and quoting happens on multiple exchanges, the ask price on one exchange may be lower than the bid price on another exchange. The literature refers to these situations as locked (ask=bid) and crossed (ask<bid) markets and they represent a market inefficiency. The BTCUSD markets we study are locked or crossed more than 99% of our sample period. This means that not only are transaction prices not integrated across exchanges but that the buy and sell orders are not integrated across exchanges. In a consolidated market where all buyers and sellers meet these orders would execute leading to transactions and higher welfare.

A major determinant of market quality is the liquidity of an asset defined as "the ease with which it is traded" (Brunnermeier and Pedersen 2008). In limit order markets such as the Bitcoin exchanges we analyze in this paper, liquidity is supplied by traders who place limit orders in the order book and is consumed by traders who submit marketable orders that execute against the limit orders in the book. While many measures of liquidity have been proposed in the empirical literature, the most direct and most widely used measures are the quoted and effective bid-ask spread (see chapter 2 in Foucault et al. 2013; Bessembinder and Venkataraman 2010; Holden, Jacobsen, and Subrahmanyam 2014). The quoted spread is the difference between the ask price and the bid price. It measures the hypothetical cost to a liquidity consumer of buying and immediately reselling an asset and is an estimate of the execution cost for small orders. The effective spread is the actual execution cost of a transaction. If the execution cost exceeds the benefit from trading, traders will refrain from trading even though the trade would have been beneficial in a hypothetical perfectly liquid market. Jennifer Huang and Wang 2010 describe the social cost to illiquidity. We find large disparities in the level of liquidity across exchanges. The quoted spreads (effective spreads) range from 0.22 (0.51) basis point (bps) at GDAX and 9.64 (7.94) bps at Bitstamp. The percentage quoted spreads on GDAX and Bitfinex are even lower than those in highly liquid equity markets where the spread is constrained by the 1-cent tick size (Rindi and Werner 2017). The price impact, an estimate of the adverse selection component of the spread defined as the difference between the midpoint between the bid and ask price directly before a transaction and the midpoint a few seconds after the transaction, is roughly the same across markets ranging from 0.76 and 1.14 bps. We also show that price disparities have decreased and liquidity has increased over our sample period.

The contrast between cryptocurrency markets and most equity markets is stark. Crypto-exchanges are almost completely unregulated yet transact roughly 3 Billion US dollars daily in BTC (this volumes includes all currency-pairs with BTC, crypto or fiat, and is not limited to BTCUSD only). The lack of regulation means, among other things, that investors on an exchange bear default and fraud risk.[1] The United States, Europe, and most developed countries have a set of financial regulations that govern the trading of securities (equities, bonds, options, and futures).[2] It is unclear if Bitcoin is a security in a legal sense and how this could affect markets.[3]

In the past most transfers of BTC has been via the decentralized and public proof-of-work Blockchain. When a participant transfers BTC to an external wallet the transfer of ownership is recorded and is observable on the public Blockchain. Exchanges allow for the transfer of Bitcoin and fiat without reporting the transfer on the Blockchain. The total volume of Bitcoin transactions therefore consists of exchange volume (off-chain) and Blockchain based transactions (on-chain). In early 2017, total exchange volume amounted to roughly 30% of total Bitcoin volume. With the beginning of the Bitcoin bull market exchanges progressively process more transfers. As of May 2018, the total exchange volume amounts to 90% of total volume (off-chain and on-chain). This means that for every Blockchain reported transfer of one Bitcoin there are exchange transfers of 9 Bitcoin at the same time, highlighting the tremendously increasing importance of exchanges. The rise of exchanges runs contrary to the philosophy of cryptocurrencies. Bitcoin was designed as an anonymous, decentralized counterpart to central bank controlled money supply. Crypto-exchanges require customers to fully reveal their identity. Customers of cryptocurrency exchanges must also transfer funds (fiat and cryptocurrency) and control over these funds to the exchange, similar to the banking system BTC was designed to circumvent.

The rest of the paper is organized as follows. Section 2 discusses the literature and Section 3 presents institutional details. Section 4 describes the data and descriptive statistics and Section 5 our results. We conclude and discuss future work in Section 6.

[1]The crypto-exchange Mt. Gox dominated the Bitcoin market, by 2013 the majority of worldwide Bitcoin volume was settled there. In spring 2014, the company declared bankruptcy, a total of 650,000 Bitcoins held in exchange accounts on the behalf of customers were lost. The recent past has also seen a number of attacks on crypto-exchanges, where hackers succeeded in illegally transferring funds from exchange-hosted user accounts.

[2]See: Reg-NMS in the US - `https://www.sec.gov/rules/final/34-51808.pdf`, or MiFID I and II in the EU - `https://ec.europa.eu/info/law/markets-financial-instruments-mifid-directive-2004-39-ec_en`; `https://ec.europa.eu/info/law/markets-financial-instruments-mifid-ii-directive-2014-65-eu_en`

[3]See: `https://goo.gl/syeCZa`

2 Literature

Our paper is related to three strands of the literature. We contribute to the literature on cryptocurrencies. Our work borrows from the market microstructure literature, in particular from the research on liquidity and market fragmentation. The market microstructure literature provides us with valuable findings on the functioning of traditional asset markets (like equity or FX markets) which we can use as a benchmark for our study on cryptocurrency markets. We also survey the literature on the role of technology and regulation in financial markets.

Much of the emerging literature on cryptocurrencies has focused on issues such as the determinants of their value (Aoyagi and Adachi 2018, Schilling and Uhlig 2018, Pagnotta and Buraschi 2018), their legal and illegal uses (Foley, Karlsen, and Putniņš 2018), their return, hedging and diversification properties (e.g. Dyhrberg 2016b, Dyhrberg 2016a, Bouri et al. 2017, Baur, Dimpfl, and Kuck 2018, Corbet et al. 2018, Brauneis and Mestel 2018a, Hu, Parlour, and Rajan 2018, Urquhart and H. Zhang 2018), or the price discovery process and the informational efficiency of Bitcoin prices (Bariviera et al. 2017, Urquhart 2016, Nadarajah and Chu 2017, Bariviera 2017). Kroeger and Sarkar 2017 and Makarov and Schoar 2018 document that Bitcoin prices at different exchanges diverge persistently. Gandal et al. 2018 provide evidence of price manipulation in Bitcoin markets. The liquidity of Bitcoin markets has not received much attention. While some papers assess liquidity based on transactions data and/or low frequency data (Brauneis and Mestel 2018b, Fink and Johann 2014, Dimpfl 2017, Shi 2018), we are aware of only one paper that uses high frequency quote data (Makarov and Schoar 2018). One of the contributions of our paper is to provide a thorough analysis of Bitcoin liquidity based on an extensive high-frequency data set.

Liquidity has been studied extensively in the literature and it is well established that illiquidity reduces asset prices and increases expected returns (e.g. Amihud and Mendelson 1986, Pastor and Stambaugh 2003, Acharya and Pedersen 2005). Changes in market architecture that reduce illiquidity are welfare-enhancing (e.g. Amihud, Mendelson, and Lauterbach 1997). Illiquidity can be caused by the existence of traders who have better information on the value of the asset than the suppliers of liquidity (Glosten and Milgrom 1985), and by other frictions including the existence of order processing costs, exchange fees, transaction taxes, and monopoly power (e.g. Glosten 1987). Corresponding to this dichotomy, bid-ask spreads can be decomposed into the price impact and the realized spread (e.g. R. Huang and Stoll 1996). While the former measures the amount that suppliers of liquidity lose to better informed traders, the latter is the gross revenue to the suppliers of liquidity and has to cover the costs associated with the other frictions listed above.

Trading in financial markets has seen revolutionary changes with technology, regulation and innovation (see ESMA 2014) reducing frictions and leading to more efficient prices and liquid markets. Trading, clearing, and settlement processes are almost completely automated, implying a substantial reduction in monitoring frictions and trading costs. Automation on NYSE and the concurrent increase in automated trading led to higher liquidity (Hendershott, Jones, and Menkveld 2011). Reducing latency on exchanges lead to more liquidity and more quote-based price discovery (Riordan and Storkenmaier 2012). Fully automated high frequency traders are associated with more price discovery (Brogaard, Hendershott, and Riordan 2014, Hasbrouck and Saar 2013, and Chaboud et al. 2014). The effect of high frequency trading on liquidity is ambiguous. Several papers conclude that high frequency trading increases liquidity (e.g. Hendershott, Jones, and Menkveld 2011, Hasbrouck and Saar 2013, Chaboud et al. 2014, Boehmer, Fong, and Wu 2015, Malinova, Park, and Riordan 2017) while others reach the opposite conclusion (e.g. Brogaard, Hendershott, and Riordan 2017, Chakrabarty, Jianning Huang, and Jain 2018). The markets we study are fully automated and are likely also populated by automated traders trading via the markets' APIs.

The Bitcoin exchanges we analyze are almost completely unregulated. In contrast, traditional financial markets have numerous rules governing the routing of market and limit orders (Reg NMS, MiFID). For instance, in the U.S. the SEC's Order Protection Rule establishes price priority across markets ensuring that an executable order executes at the market that is offering the best price (Macey and O'Hara 1997, Hansch, Naik, and Viswanathan 1999). This rule and similar rules elsewhere limit the likelihood of trades executing at prices inferior to those available at other trading venues. The rules thus make it less risky for traders to route their orders to alternative trading venues and thus contribute to an increase in the fragmentation of trading across venues. In the EU under MiFID the task of assuring best execution of customer orders in the presence of fragmented markets lies with the brokers (Degryse 2009). The evidence is mixed on fragmentation with some papers suggesting that fragmentation is positive for liquidity (O'Hara and Ye 2011) while other papers suggest that fragmentation is good only for large firms and bad for small firms (Haslag and Ringgenberg 2016). Several studies on European stock markets (Degryse, Jong, and Kervel 2015, Gresse 2017, Hengelbrock and Theissen 2009) document predominantly positive effects of increased competition between trading venues. The rules certainly ensure that small investors generally receive the best price for their orders and that the prices across trading venues are integrated. No such regulation exists in the case of cryptocurrency markets, making them an ideal subject to investigate the developments of decentralized and unregulated markets.

3 Institutional Details

We use data from three of the largest cryptocurrency spot trading platforms that trade Bitcoin against the US Dollar (BTCUSD): Bitfinex, Bitstamp and GDAX. All platforms operate an electronic central limit order book with orders being matched based on price-time-priority. Along with Bitcoin's price rally over the last few years, a vast number of new trading platforms and exchanges emerged (as of June 15 2018, coinmarketcap.com lists over 11,000 cryptocurrency markets). One would expect the law of one price to enforce the price of Bitcoin to be equal across all exchanges. However, transaction data for our three markets points at price differences violating the law of one price. Figure 1 plots the maximum percentage difference of hourly, value weighted prices for these markets for the period of January 2015 to May 2018.

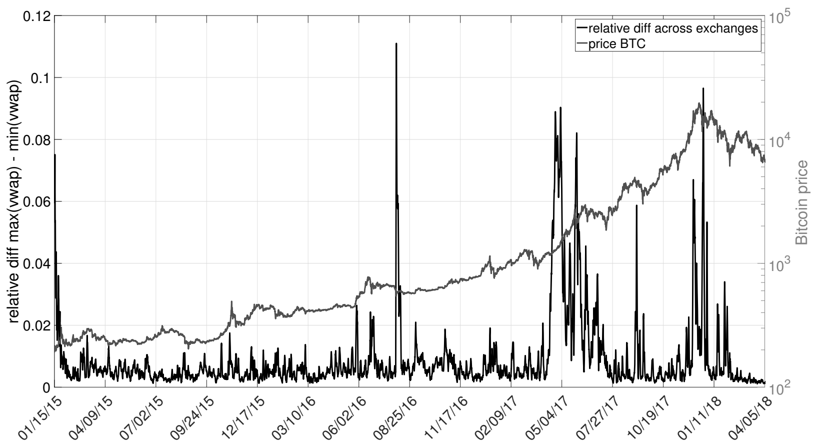

Figure 1: Percentage price difference of hourly value weighted average prices (vwap) for Bitfinex, Bitstamp and GDAX. Bitfinex data is taken from Bitfinex' website whereas time series of hourly prices of Bitstamp and GDAX are derived from a full record of all transactions provided by bitcoincharts.com.

Figure 1 shows price differences of up to 11%. Beginning in February 2018, and after Bitcoin's price peaked at its all-time-high in December 2017, price differences decrease (see right hand end of the time series plotted in Figure 1). We study this period in more detail in the following sections.

The tremendous increase in the Bitcoin price in 2017 goes hand in hand with a sharp increase in exchange traded volume. Bitcoin transactions may either be conducted on-chain via the Blockchain or off-chain on cryptocurrency exchanges. While data on aggregated exchange volume is readily available (e.g. on coinmarketcap.com), Blockchain volume may only be estimated. This is due to the fact that the Bitcoin protocol requires users to fully use their funds in every transaction.[4] We use data from blockchain.info to estimate on-chain transactions. Figure 2 shows that the proportion of these transactions considerably fluctuated at around 30% from 2014 until the beginning of 2017. Since then, off-chain volume on crypto-exchanges has taken the leading part with a meanwhile stable proportion of almost 90%.

[4]If, for example, a user's balance amounts to 10 BTC, a transaction of 1 BTC to another address in the network consists of one transaction with a value of 10 BTC and another transaction with a value of 9 BTC, returned to the sender (but possibly to a different wallet of the senders'; the returned amount is referred to as change). The gross volume thus is 10 BTC, effectively, only 1 BTC changes hands.

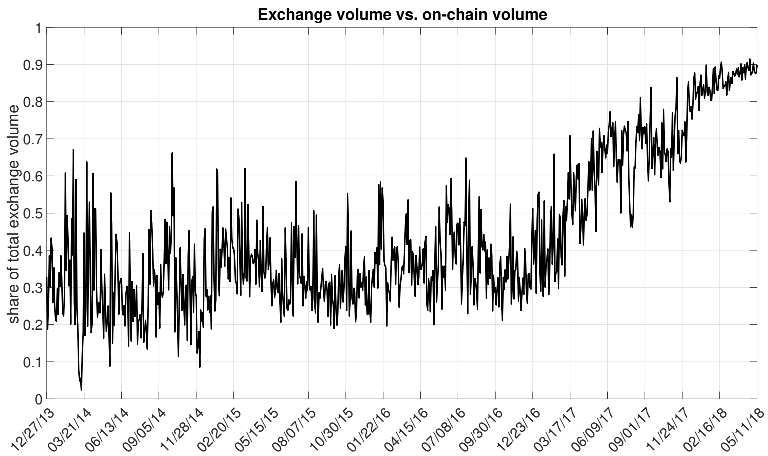

Figure 2: Share of exchange transferred Bitcoins to total Bitcoin volume defined as exchange volume plus Blockchain volume. Aggregate exchange volume stems from coinmarketcap.com, (estimates of) on-chain volume is taken from blockchain.info.

As of June 15, 2018, Hong Kong-domiciled Bitfinex (www.bitfinex.com) was the world's largest Bitcoin trading platform by volume against the US Dollar. Orders at this exchange are executed under a maker/taker fee schedule[5] with fees depending on the executed USD equivalent volume over a 30-day interval. The minimum order size is adjusted periodically and is intended to range between 10 and 25 USD equivalent value. The minimum tick size for BTCUSD is 0.1 USD. Bitfinex also offers a peer-to-peer financing functionality by means of a separate financing order book allowing for a maximum total bitcoin-to-equity ratio of 3 1/3 to 1. Additionally, investors can borrow digital tokens for short sales via the financing order book. Bitfinex also operates an OTC trading desk allowing traders to make larger trades directly with a counterparty outside the public order books. Bitfinex is currently not doing business with permanent US residents.

Bitstamp (www.bitstamp.net) is the most well-known European cryptocurrency trading platform, based in Luxembourg. Unlike Bitfinex, Bitstamp charges a uniform trading fee, i.e. liquidity makers and takers pay the same execution fee. The exchange stipulates a minimum trade size of 5 USD and the minimum tick size for Bitcoin traded against USD is 0.01 USD

[5] A transaction occurs when a marketable order hits one (or several) limit orders sitting in the limit order book. The marketable order "takes" the liquidity that was offered (or "made") by the limit order. Exchanges charge traders execution fees which traditionally were identical for the "maker" and "taker" orders. Under a maker/taker fee schedule these fees are different. The most common case is a maker fee that is lower than the taker fee, possibly even negative (i.e. a rebate rather than a fee). Low or negative maker fees provide incentives to limit order traders to quote a narrow bid-ask spread. For a detailed analysis of maker/taker fee schedules see Colliard and Foucault 2012 and Malinova and Park 2015.

Like Bitfinex and Bistamp, US-based digital asset exchange GDAX (www.gdax.com) charges a taker fee based on the individual USD trading volume executed in the last 30 days. Unlike the two other platforms, however, GDAX does not impose execution fees on limit orders, i.e. the maker fee is 0%. The minimum BTCUSD order size equals 0.001 BTC, the minimum tick size on GDAX is 0.01 USD.

Table 1 lists comprehensive details on institutional aspects of all three exchanges under consideration in this paper (as of June 15, 2018).

Table 1: Margin trading rules, fee structures for liquidity provider and liquidity taker, tick sizes for the three exchanges. All information presented in this table stems from the exchanges' websites as of June 15, 2018.

Property	Bitfinex	Bitstamp	GDAX
Margin Trading	Peer-to-peer margin funding platform, allowing for long positions with up to 3.3x leverage as well as short sales	Announced, but not realized yet	Abolished as of June 2017
Maker Fee	Depending on the individual USD trading volume executed in the last 30 days: $>$ 0: 0.10% $\geq$ 500,000: 0.08% $\geq$ 1,000,000: 0.06% $\geq$ 2,500,000: 0.04% $\geq$ 5,000,000: 0.02% $\geq$ 7,500,000: 0.00%	Depending on the individual USD trading volume executed in the last 30 days: $>$ 0: 0.25% $\geq$ 20,000: 0.24% $\geq$ 100,000: 0.22% $\geq$ 200,000: 0.20% $\geq$ 400,000: 0.15% $\geq$ 600,000: 0.14% $\geq$ 1,000,000: 0.13% $\geq$ 2,000,000: 0.12% $\geq$ 4,000,000: 0.11% $\geq$ 20,000,000: 0.10%	0% irrespective of the individual trading volume
Taker Fee	Depending on the individual USD trading volume executed in the last 30 days: $>$ 0: 0.20% $\geq$ 10,000,000: 0.18% $\geq$ 15,000,000: 0.16% $\geq$ 20,000,000: 0.14% $\geq$ 25,000,000: 0.12% $\geq$ 30,000,000: 0.10%	Depending on the individual USD trading volume executed in the last 30 days: $>$ 0: 0.25% $\geq$ 20,000: 0.24% $\geq$ 100,000: 0.22% $\geq$ 200,000: 0.20% $\geq$ 400,000: 0.15% $\geq$ 600,000: 0.14% $\geq$ 1,000,000: 0.13% $\geq$ 2,000,000: 0.12% $\geq$ 4,000,000: 0.11% $\geq$ 20,000,000: 0.10%	Depending on the individual USD trading volume executed in the last 30 days: $>$ 0: 0.30% $\geq$ 10,000,000: 0.20% $\geq$ 100,000,000: 0.10%
Min Tick Size	0.10 USD	0.01 USD	0.01 USD

Empirical market microstructure literature has established evidence on how minimum tick sizes and the structure of execution fees affect liquidity on equity markets. The minimum tick size imposes a lower lower bound on the quoted bid-ask spread. Several studies document a positive relation between the tick size, spreads and depth (e.g. Harris 1999, Rindi and Werner 2017). With 0.10 USD, Bitfinex has adopted a minimum tick size ten times larger than those at Bitstamp and GDAX. We therefore expect Bitfinex to have higher spreads and higher depths.

Rational traders will take into account the execution fee when deciding upon the prices at which they quote and trade (e.g. Colliard and Foucault 2012, Malinova and Park 2015). Consequently, a lower maker fee will result in a lower quoted spread. By this logic we expect GDAX (which charges no maker fee at all) to have the lowest quoted spreads. Bitfinex has lower maker fees than Bitstamp. However, because Bitfinex also has a much larger minimum tick size no clear prediction with respect to the relative size of the spread on these two exchanges can be made.

Bitfinex offers margin trading via its financing order book. This feature allows traders to sell short and to open leveraged long positions. Several empirical studies conducted using equity market data (e.g. Boehmer, Jones, and X. Zhang 2013, Beber and Pagano 2013) suggest that the possibility to short sell positively affects trading activity and liquidity. To the extent that this result carries over to Bitcoin trading we expect higher trading volume and higher liquidity on Bitfinex.

4 Data and Methodology

We compiled a high-frequency data set that covers the six-months period from 12/15/2017 06:00 UTC to 06/15/2018 06:00 UTC, a total of $T = 262,080$ minutes. Over this period we used Matlab to continuously access the public and freely accessible REST API of Bitfinex, Bitstamp and GDAX. They provide live information on transactions and the current state of the order book. All public endpoints at each of these exchanges use GET requests for different types of information. We basically request records on 'Trades' / 'Transactions' and the 'Orderbook'. Depending on the exchange, request parameters vary. For instance, Bitstamp only provides the full order book (with usually thousands of entries) whereas order book requests at Bitfinex and GDAX may be limited to the 50 best orders on each side of the market.

We performed a brief test of server response times (the time elapsed between the API call and the server's response). For a sample of 1,000 requests Bitfinex on average responds within 0.34 sec (0.33 sec) on a transactions (order book) request. The corresponding values for Bitstamp (GDAX) are 0.22 sec and 0.49 sec. (0.88 sec and 0.91 sec) for transaction and order book requests, respectively.

4.1 Transaction and Order Book Data

We record an overall number of transactions during our sample period of over 40 million. For each transaction the data set includes the price and the corresponding trading volume, a UNIX time stamp, a unique exchange-specific ID and a trade indicator which indicates whether a transaction was buyer-initiated or seller-initiated. The upper part of Table 2 provides a short description of this data. The total number of transactions ranges from 7.65 million (corresponding to one transaction every 2.06 seconds) on Bitstamp over 13.43 million transactions (one transaction every 1.17 seconds) for GDAX to 19.23 million transactions (one transaction every 0.82 seconds) for Bitfinex. The fraction of buyer-initiated trades is less than 50% (at 48.51%) on Bitfinex and is above 50% (at 57.10% and 54.28%, respectively) on Bitstamp and GDAX.

We observe several time intervals with gaps in the data (denoted *outages* in the sequel). These may be due to actually missing trading activity, technical problems (failure of the internet connection, no response from the server etc.), or exchange-specific trading halts (e.g. due to maintenance, updates or hacker attacks). We identify between 2,408 (GDAX) and 5,106 (Bitstamp) intervals exceeding 60 seconds (1 minute), between 504 and 706 intervals exceeding 600 seconds (10 minutes) and between 79 and 87 intervals exceeding 1,800 seconds (30 minutes) without transaction data.

Table 2: Description of transactions and order book data for each exchange. The upper part reports the total number of transactions (# TA), the average time between two transactions (avg interim time), and the percentage of total transactions which are buyer initiated according to the trade indicator variable provided by the respective exchange. The lower part shows the number of order book snapshots (# OB snapshots) and the average time between two order book snapshots (avg interim time). Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

	Bitfinex	Bitstamp	GDAX	Overall
# TA	19,226,158	7,647,337	13,434,845	40,308,340
avg interim time	0.8181s	2.0567s	1.1707s	
Trade indicator Buy	48.51%	57.10%	54.28%	
# OB snapshots	1,711,356	1,723,561	1,747,979	5,182,896
avg interim time	9.1904s	9.1253s	8.9978s	

Besides transactions we retrieve order book data from the three trading platforms. Specifically, we collect the 50 best bid and best ask prices with corresponding volumes. As for transaction data the frequency at which this information is available depends on server response times and the speed of the internet connection. As can be seen from the lower panel of Table 2 we collect between 1.7 million (Bitfinex) and 1.75 million (GDAX) UNIX time-stamped order book snapshots. On average, across the three exchanges, order book data is obtained every 9 seconds. As for transaction data we observe a considerable number of outages. There are between 3,814 (GDAX) and 7,510 (Bitfinex) intervals exceeding 60 seconds without an order book snapshot. The numbers of intervals without order book snapshots exceeding 600 seconds and 1,800 seconds are roughly equal across the three exchanges and amount to approximately 810 and 210, respectively.

The 50 orders that each snapshot contains represent considerable volume. For Bitfinex the mean total bid (ask) dollar volume amounts to 1.37 mio USD (1.38 mio USD), for Bitstamp the 50 best bid orders (ask orders) on average sum up to 1.57 mio USD (1.39 mio USD) and for GDAX the mean total dollar volume in bid (ask) orders amounts to 0.54 mio USD (0.55 mio USD).

We sort all data into one-minute intervals. We use the following notation.

$P^b_{t,i}$ ($P^a_{t,i}$) refers to the best bid (best ask) of order book snapshot i in interval t, and $m_{t,i}$ denotes the quote midpoint. For each one-minute interval $t \in \{1, ..., T\}$ we compute liquidity measures based on the first order book snapshot observed in that interval (i.e. we set i equal to 1). Since some of our liquidity measures (such as the price impact) also require data on subsequent transactions, the first order book of the interval is the best choice. $P_{t,j}$ denotes the price of transaction j in interval t.

We include in our analysis only intervals for which we observe at least one sequence consisting of a first order book, one subsequent transaction and a second order book thereafter. Applying this filter rule to the total number of 262,080 one-minute intervals results in 231,943 / 232,660 / 233,575 one-minute intervals for Bitfinex / Bitstamp / GDAX. We refer to this data set as the exchange-specific data set because it contains, for each exchange, all one-minute intervals with complete information for that exchange. We further compile a synchronized, or synced, data set that only contains those one-minute intervals for which complete information is available for all three exchanges. This data set contains 224,424 one-minute intervals.

4.2 Outages

As documented above there is a considerable number of outages, i.e. intervals with missing data. There are more instances of missing order book data than instances of missing transaction data. This is because, with every request, we obtain the transactions of the previous minute (Bitstamp) or the most recent 100 transactions (Bitfinex and GDAX). Consequently, only no-response periods of more than two minutes will result in a loss of transaction data. The same is not true for order book data. Here, a request only delivers the current state of the order book. We compare the time-stamp of each order book snapshot to the previous snapshot and discard duplicates. Outages may occur for technical reasons, because of software updates (e.g. on April 11, 2018 [Bitfinex]), maintenance activities (March 21, 2018 [Bitfinex], June 4, 2018 [GDAX]) or DDoS attacks (e.g. Dec 17, 2017, Dec 31, 2017 and June 5, 2018 [Bitfinex]). Further, outages may occur randomly or systematically. If, for example, the probability of a no-response event increases in the system load, the failure event is systematic. It is important to know whether outages occur systematically for at least two reasons. First, if they did, our data set would potentially be affected by a selection bias because the probability of recording an observation would be lower in times of high system load (e.g. times of stress). Second, and more importantly, if outages occurred systematically traders' access to the market would be limited in times of high system load. However, a market that is less accessible in times of stress is less attractive.

We address this issue by analyzing market behavior around outages. When data for all three exchanges is missing simultaneously we assume that our infrastructure is the source of the failure. When data for one exchange is missing and data for the other exchanges is available, we assume that an external reason causes the outage. We only analyze the latter cases and correspondingly define an *outage* as an interval for which exactly one of the three exchanges fails to report the current state of the order book via the API.[6] If an exchange delivers no order book data traders must either refrain from trading, or have to base their trading decisions on stale information.

Table 3: Server disconnectivity. Analysis of transactions, dollar volume and price volatility for server malfunctions while the other two exchanges' APIs are responding. The upper panel reports mean and median results for the number of transactions, dollar volume and volatility based on the synced dataset comprising 224,424 on-minute intervals. The middle part reports results on all one-minute intervals where one server is down. The lower part reports results on server outages where one serve is not responding for at least 5 consecutive one-minute intervals. Numbers refer to the mean, the median is reported in parentheses. Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

	Bitfinex	Bitstamp	GDAX
synced data $N = 224,424$ intervals, numbers report mean (median)			
numTA	82.41 (57)	32.61 (22)	56.93 (33)
DV [1,000 USD]	303.71 (145.04)	102.30 (38.66)	119.80 (49.35)
Vola [bp]	10.54 (5.73)	11.51 (7.29)	8.41 (2.76)
Analysis of server disconnectivity I			
Server 'down'	1,708	447	698
Bitfinex down			
num TA	-	28.08 (19.5)	45.05 (26)
DV [1,000 USD]	-	83.30 (31.79)	88.79 (35.09)
Vola [bp]	-	9.31 (6.03)	6.37 (1.15)
Bitstamp down			
num TA	82.45 (36)	-	56.57 (22)
DV [1,000 USD]	339.78 (69.53)	-	136.15 (28.74)
Vola [bp]	12.64 (7.20)	-	11.64 (5.14)
GDAX down			
num TA	58.38 (36.5)	26.91 (15.5)	-
DV [1,000 USD]	208.62 (83.96)	82.75 (21.12)	-
Vola [bp]	8.63 (3.69)	11.01 (6.35)	-
Analysis of server disconnectivity II			
Server $\geq$ 5min 'down'	458	128	345
Bitfinex $\geq$ 5min down			
num TA	-	23.03 (16)	24.73 (14)
DV [1,000 USD]	-	57.41 (21.23)	29.32 (12.08)
Vola [bp]	-	5.45 (4.03)	2.53 (0.01)
Bitstamp $\geq$ 5min down			
num TA	73.03 (58.5)	-	50.58 (31)
DV [1,000 USD]	231.64 (129.57)	-	125.36 (56.65)
Vola [bp]	9.64 (6.42)	-	7.57 (2.95)
GDAX $\geq$ 5min down			
num TA	40.93 (27)	21.97 (13)	-
DV [1,000 USD]	134.42 (54.38)	65.18 (15.52)	-
Vola [bp]	7.09 (2.25)	9.43 (5.13)	-

The results of our analysis are shown in Table 3. The upper panel shows the unconditional mean and median of the number of transactions, the dollar trading volume and the volatility (measured by the squared one-minute return)

[6] Note that transaction data may be available for these intervals because, as noted above, a successful data request delivers the transactions of the most recent two minutes.

for each exchange. These figures serve as benchmark values and are based on the synchronized data set.

The middle panel considers cases in which one exchange is not responding. We count 1,708 [447; 698] intervals in which Bitfinex [Bitstamp; GDAX] is not responding while the other exchanges are. Bitfinex thus is the exchange with the highest outage rate while Bitstamp has the lowest rate.[7]

We analyze the contemporaneous activity on those two exchanges that are unaffected by the outage. Consider the line labelled "Bitfinex down". While Bitfinex is down, Bitstamp (GDAX) records an average of 28.1 (45.1) transactions per minute, slightly less than the unconditional mean of 32.6 (56.9) shown in the upper panel of the table. Similarly, the Bitstamp and GDAX dollar volume and volatility are also below their unconditional means when Bitfinex is non-responding. We obtain similar results when GDAX is not responding (see the line labelled "GDAX down"). Trading activity and volatility at Bitfinex and Bitstamp are slightly below their unconditional means. These results imply that Bitfinex and GDAX are non-responding during intervals of (slightly) below-average trading intensity and volatility. Such a pattern may arise when the outages are due to maintenance or software updates which are deliberately timed to take place at times of low activity.

We obtain different results when Bitstamp is non-responding (see the line labelled "Bitstamp down"). In this case, the dollar volume on Bitfinex and GDAX is more than 10% and volatility is more than 20% above their unconditional means. This pattern suggests that the response probability of Bitstamp is lower during times of high trading volume and volatility.

We repeat the analysis considering only outages that last at least 5 minutes. Results are shown in the lower panel of Table 3. There are 10 [18; 31] of these outages covering a total of 458 [128; 345] minutes for Bitfinex [Bitstamp; GDAX]. We consistently find that trading activity as well as volatility are below their unconditional means during the outages. The most likely reason is, again, that these outages occur because of scheduled events such as maintenance or software updates.

In summary, we find that Bitfinex has the highest non-response rate, but outages at this exchange, as at GDAX, appear to not occur at times of high system load. On the other hand, Bitstamp has the lowest overall non-response ratio, but there is some evidence that this exchange is non-responding in a systematic way, i.e. in times of market stress. We further find that longer outages (i.e. those exceeding 5 minutes) are most likely due to scheduled events taking place at times of low activity.

[7] We note that the outage rate may depend on the geographical distance between our server (based in Austria) and the exchange server. In this case we would expect higher outage rates for Bitfinex and GDAX than for Bitstamp. However, outages due to geographic distance should be unsystematic.

4.3 Liquidity Measures

Based on the one-minute interval data, we calculate the following standard measures of liquidity:

1. Percentage Quoted Spread QS. For interval t, this measure is defined as $QS_t = \frac{P^a_{t,1} - P^b_{t,1}}{m_{t,1}}$

2. Percentage Effective Spread ES. For interval t, the effective spread is defined as $ES_t = 2 \cdot Q_{t,j'} \cdot \frac{P_{t,j'} - m_{t,1}}{m_{t,1}}$, where j' refers to the first transaction after the order book snapshot was recorded and $Q_{t,j'}$ is a trade indicator variable ($Q_{t,j'} = 1$ for a buyer-initiated trade, $Q_{t,j'} = -1$ for a seller-initiated trade).

3. Percentage Price Impact PI. We consider the first transaction that occurs after the order book snapshot and record the sign of the transaction. We then consider the quote midpoint $m_{t,i+1}$ from the next order book snapshot. The price impact is then calculated as $PI_t = \frac{Q_{t,j'} \cdot (m_{t,i+1} - m_{t,i})}{m_{t,i}}$.

4. Average BBO Depth $AvgD$. Depth is defined as $AvgD_t = \left(P^a_{t,1} \cdot V^a_{t,1} + P^b_{t,1} \cdot V^b_{t,1}\right)/2$, where $V^{a,(b)}_{t,1}$ denotes the volume associated with the best ask and bid price, respectively.

5. Dollar Volume DV. For interval t, volume is defined as $DV_t = \sum_j P_{t,j} \cdot V_{t,j}$, where $V_{t,j}$ is the amount of Bitcoins traded in transaction j.

6. Number of transactions $numTA$. This measure is defined as the total number of individual transaction observed in interval t.

7. Order Imbalance OI. For interval t, order imbalance is defined as $OI_t = \frac{\sum_{j,Q_{t,j}=1} Q_{t,j} - \sum_{j,Q_{t,j}=-1} |Q_{t,j}|}{\sum_j |Q_{t,j}|}$.

8. Order Imbalance Volume OIV. For interval t, this measure is defined as $OIV_t = \frac{\sum_{j,Q_{t,j}=1} P_{t,j} \cdot V_{t,j} - \sum_{j,Q_{t,j}=-1} P_{t,j} \cdot V_{t,j}}{\sum_j P_{t,j} \cdot V_{t,j}}$

The quoted bid-ask spread is a valid measure of execution costs only for small trades, i.e. trades the size of which does not exceed the depth available at the best quotes. Larger market orders will walk up or down the book and will thus partly execute at worse prices. To assess the liquidity of the Bitcoin markets for larger trades we use the order book data to calculate the weighted average price at which a buy and a sell order of a given size Y would execute. The weighted average price WAP in interval t for executing a transaction of size Y USD given the current order book is defined as $\frac{\sum_j A_j \cdot V_j}{\sum_{j=1}^J V_j}$ subject to $\sum_{j=1}^J A_j \cdot V_j = Y$ where A_j denotes the j^{th} order in the order book, and A_j, V_j are the price and volume of the j^{th} order, respectively. Note that the J^{th} order may be subject to partial execution, depending on the outstanding dollar volume in order to entirely fill the transaction volume Y. We calculate weighted average prices for different order sizes. Specifically, we set Y equal to 500, 2,000, 40,000, and 100,000 USD. These values roughly correspond to the median, the third quartile, the 99% and the 99.9% quantiles of the (aggregated) trade size distribution.

5 Results

We present our results in three steps. We first present evidence on the trading activity on the Bitcoin exchanges and document its intra-day and intra-week patterns. This is important because the exchanges operate 24 hours a day 365 days a year. Thus, in contrast to equity markets, there are no trading halts overnight or over the weekend. We then return to the issue of market integration and analyze the frequency and magnitude of market crosses, i.e. situations in which the bid price in one market exceeds the ask price in another market. This is important because, absent any other frictions, crosses constitute arbitrage opportunities and are thus evidence against market integration. Finally, we analyze the liquidity of the Bitcoin exchanges. This is important for at least two reasons. First, while liquidity is usually considered to be an important determinant of market quality, little is known on the liquidity of Bitcoin exchanges. Second, if markets are integrated, liquidity on these markets should move in lockstep. Therefore, an analysis of liquidity also provides insights into the degree of market integration.

5.1 Trading

We start with results on trading activity on our three cryptocurrency exchanges. Using individual platform data, the upper part of Table 4 shows an average one-minute number of transactions ranging from 32.59 (Bitstamp) to 81.85 (Bitfinex). The average dollar volume for single transactions is lowest on GDAX (2,070 USD) and highest on Bitfinex (3,648 USD).[8] The trade size distribution is heavily skewed, as seen in the large difference between the mean and median volumes.

Single transactions also serve as a basis for estimates on one-minute interval price volatility. In particular, we use the quantity: $\sqrt{[\ln(P_{t,1}/P_{t-1,1})]^2}$ based on the first observed transactions in each interval. Note that we normalize this volatility measure by the factor $\sqrt{(60/time-gap)}$, such that irrespective of the actual time elapsed between two transaction we obtain an estimate of the sixty-second-volatility. We find volatility to be highest on Bitstamp (11.49 bp) and lowest on GDAX (8.39 bp).

[8] We note that some trades observed in our dataset are not in line with the minimum order size rules reported on the exchanges' websites. This is most likely partially executions being reported as multiple transactions in the data. We observe several transactions with a size of 1 Satoshi (1/100mio BTC).

The lower part of Table 4 displays results for trading activity and volatility based on the synchronized dataset. They are almost identical to those obtained from the exchange-specific data sets.

Table 4: Descriptive statistics for one-minute intervals for number of transactions, dollar volume and volatility. The upper panel reports results for exchanges' individual data with a varying number of intervals with full data availability.'no data [1 min]' refers to the total number of one-minute intervals, adjacent or not, without data. 'no data [adjacent intervals]' reports the number of intervals of varying length (1,2, ... minutes) without data. The lower panel reports results for the synced data set with full availability of data for all three exchanges. Numbers refer to the mean except for total num intervals, the median is reported in parentheses. Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

	Bitfinex	Bitstamp	GDAX
individual exchange data, numbers report mean (median)			
total num intervals	231,943	232,660	233,575
numTA	81.85 (57)	32.59 (22)	56.68 (33)
DV [1,000 USD]	301.43 (143.11)	102.28 (38.31)	119.43 (48.69)
Vola [bp]	10.50 (5.69)	11.49 (7.27)	8.39 (2.67)
no data [1 min]	30,137	29,420	28,505
no data [adjacent intervals]	8,971	7,913	5,795
synced data $N = 224,424$ intervals, numbers report mean (median)			
numTA	82.41 (57)	32.61 (22)	56.93 (33)
DV [1,000 USD]	303.71 (145.04)	102.30 (38.66)	119.80 (49.35)
Vola [bp]	10.54 (5.73)	11.51 (7.29)	8.41 (2.76)

Figures 3 and 4 show how trading activity and volatility evolve over the days of the week and over the time of the day. The number of transactions, the USD dollar volume, and volatility all peak towards the end of the week and are lowest on Saturdays and Sundays. Daily activity peaks at roughly 15:00 to 16:00 UTC. This coincides with the market opening on NYSE and NASDAQ in US equity markets.

5.2 Crossed Markets

In the absence of frictions we would expect that the highest bid price of any market (the consolidated best bid) is lower than the lowest ask price of any market (the consolidated best ask). We refer to this situation as a "normal" market. A situation in which the consolidated best bid is exactly equal to the consolidated best ask is known as a locked market, or a lock. Finally, a situation in which the highest bid exceeds the lowest ask is known as a crossed market or, in short, a cross. If there were no transaction costs other than the bid-ask spread, a crossed market would imply an arbitrage opportunity.

For every one-minute interval in the synchronized data set we check whether the market is locked or crossed. We repeat the analysis considering pairs of two markets (i.e. Bitfinex and Bitstamp, Bitfinex and GDAX, and Bitstamp and GDAX). The results are shown in Table 5. They imply that the Bitcoin

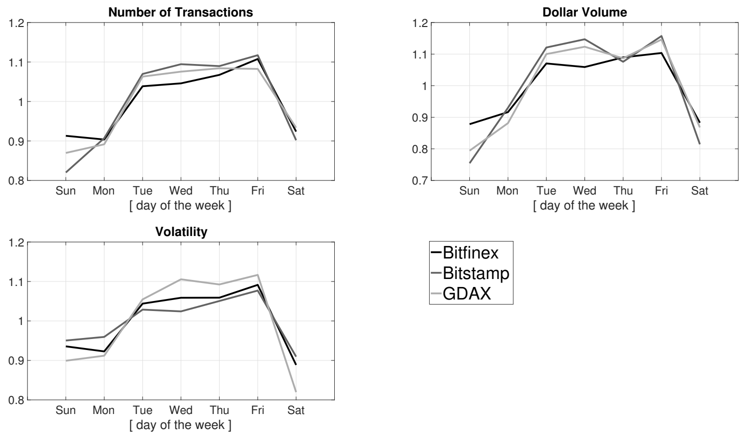

Figure 3: Temporal patterns over the course of one week, covering Sunday to Saturday. All subplots depict relative deviations from the overall mean. The upper left subplot shows the number of transactions, the upper right subplot plots the dollar volume, the lower plot refers to volatility respectively. Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

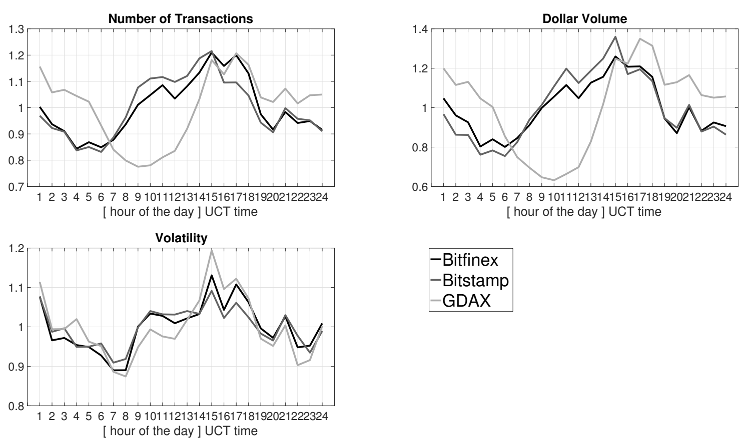

Figure 4: Temporal patterns over the course of one day, covering hours 1 to 24. All subplots depict relative deviations from the overall mean. The upper left subplot shows the number of transactions, the upper right subplot plots the dollar volume, the lower plot refers to volatility respectively. Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

market is crossed most of the time. When considering all three exchanges jointly, there are only 1,350 (out of a total of 224,424) one-minute intervals in which the market is normal. In 290 intervals the market is locked while in 222,784, or 99.27%, of the intervals the market is crossed. On average across all observations the highest bid price exceeds the lowest ask price by 61.87 USD, or 50.42 basis points.

The pairwise comparisons reveal that Bitstamp and GDAX is the most integrated pair of markets (168,537 crosses with an average difference between highest bid and lowest ask of 20.3 bp) while Bitfinex and GDAX is the least well integrated pair of markets (218,964 crosses with an average bid-ask difference of 46.3 bp).

Table 5: Number of one-minute intervals with normal / locked / crossed markets based on the synced dataset with 224,424 intervals with full data availability. BF, BS and GX refer to Bitfinex, Bitstamp and GDAX respectively. Columns avg abs (rel) spread report the mean of $\max[P^b_{it}] - \min[P^a_{it}]$ and $(\max[P^b_{it}] - \min[P^a_{it}])/(.5 \cdot \max[P^b_{it}] + .5 \cdot \min[P^a_{it}])$ for all intervals irrespective of the condition of the markets. A positive spread indicates crossed markets. Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

Set of markets	normal	locked	crossed	avg abs spread	avg rel spread
All markets	1,350	290	222,784	61.87 USD	50.42 bp
BF and BS	45,781	554	178,089	35.26 USD	29.59 bp
BF and GX	4,620	840	218,964	57.12 USD	46.25 bp
BS and GX	54,889	998	168,537	26.20 USD	20.30 bp

Figure 5 shows how the fraction of crossed markets, measured over 72-hours intervals, evolves over time. When all three markets are considered jointly, no trend is apparent. The percentage of crossed markets fluctuates between 95% and 100%. The same is true when bid and ask quotes from Bitfinex and GDAX are considered. The two other pairs (Bitstamp and GDAX; Bitfinex and Bitstamp) appear to become more integrated over time. The percentage of crosses, although very volatile, decreases over time.

As noted previously, frictions such as capital transfer controls, transaction costs and the delay in confirming Bitcoin transfers impede arbitrage activities. Thus, the existence of a large percentage of crossed markets in itself is not surprising. We therefore also consider the magnitude of the crosses, defined as the difference between the highest bid and the lowest ask. Figure 6 shows these differences, separately for each pair of exchanges. Each sub-figure contains two graphs, one for those intervals in which the first exchange's bid exceeds the second exchange's ask and one for the reverse case.

The figures clearly imply that the differences between highest bids and lowest asks have decreased considerably over time. We thus conclude that markets have become more integrated.

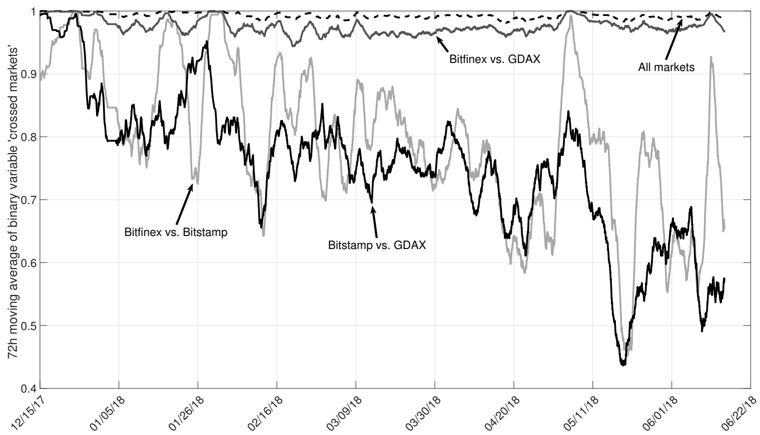

Figure 5: This figure shows a 4,320 minute (72 hours) moving average of the binary variable 'crossed markets' over time. The dashed line refers to the set of all three market, whereas solid lines refer to pairwise comparisons. The light gray (gray) [black] line shows data for Bitfinex vs. Bitstamp (Bitfinex vs. GDAX) [Bitstamp vs. GDAX]. Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

5.3 Liquidity

Thus far we have considered market integration at the level of transaction prices and price quotations. We now turn to another important dimension, namely, liquidity. In the absence of frictions one would expect that differences in transaction costs between trading venues are either non-existent, or reflect other relative advantages or disadvantages of the respective trading venues. To assess whether, or to which extent, this is the case, we now turn to an analysis of the liquidity of the three trading venues under consideration.

Table 6 reports descriptive statistics for all liquidity measures outlined in section 4.3. The measures were derived from the synced data set. GDAX has by far the lowest quoted and effective bid-ask spreads (mean 0.22 bp and 0.51 bp, respectively). This result comes as no surprise because, as outlined in section 3, GDAX has a very low minimum tick size (0.01 USD) and the most attractive fee structure for liquidity providers (maker fee = 0%).

Quoted and effective spreads on Bitfinex are only moderately higher than those on GDAX, at 1.24 bp and 1.09 bp, respectively. Spreads on Bitstamp are larger by an order of magnitude, at 9.64 bp (quoted spread) and 7.94 bp (effective spread). The two upper panels of Figure 7 plot the quoted and effective spread, separately for each exchange, over time. The figure confirms the ranking of the three exchanges (i.e. spreads are lowest on GDAX and highest on Bitstamp), and shows that this ranking is stable throughout the sample period. The figure further reveals that both quoted and effective spreads have decreased

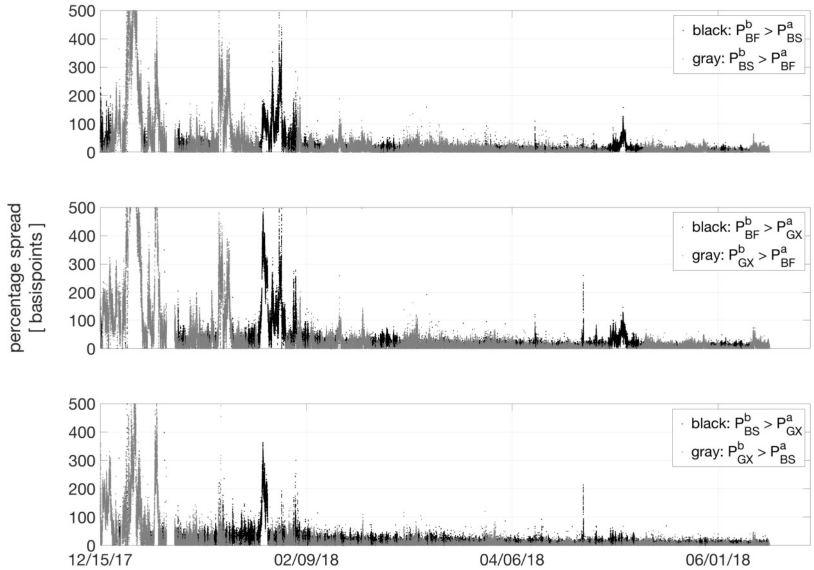

Figure 6: This figure shows 3 subplots for the percentage spread in intervals of crossed markets for pairs of 2 out of 3 exchanges respectively. The upper subplot refers to Bitfinex (BF) and Bitstamp (BS), the middle subplot presents results for Bitfinex and GDAX (GX) and the lower subplot depict data from the pair Bitstamp and GDAX. Black and gray dots show results depending on which exchange's bid exceeds the other exchange's ask quote. Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

over time.

The trading protocols of Bitfinex and Bitstamp differ on at least three dimensions. Bitfinex has a higher minimum tick size (which should result in higher spreads) but also has lower maker fees and offers margin trading and short selling opportunities (both of which arguably increase liquidity and thus should result in lower spreads). Apparently the effects of the latter features overcompensate the effect of the larger minimum tick size.[9] This finding is consistent with evidence from equity markets suggesting that the possibility to short sell positively affects liquidity (e.g. Boehmer, Jones, and X. Zhang 2013, Beber and Pagano 2013).

In section 4.2 we document that the probability of not getting a response to data requests at Bitstamp is higher at times of high trading volume and volatility. Consequently, the suppliers of liquidity on Bitstamp face an elevated risk

[9]We cannot rule out, though, that other institutional features which we have not identified cause the differences in spreads.

of not being able to continuously monitor the market particularly in those situations where this would be most important. A reasonable response to this risk is to increase quoted bid-ask spreads by an amount that compensates liquidity suppliers for the risk. Such a compensation will contribute to higher spread levels on Bitstamp.

The price impact is an estimate of the adverse selection component of the spread, i.e. of the amount that suppliers of liquidity lose to informed traders. Depicted in the lower left panel of Figure 7 we find that the price impact is lowest on Bitfinex (mean 0.76 bp), followed by GDAX (0.9 bp) and Bitstamp (1.14 bp). These values are much closer to each other than those for the quoted and effective spreads presented above. We therefore conclude that differences in adverse selection costs do not explain the spread differences between the exchanges.

The difference between the effective half-spread and the price impact (known as the *realized spread*) is an estimate of the gross revenue of the suppliers of liquidity. Realized spreads are negative on GDAX (at 0.51/2 - 0.9 = -0.65 bp), also negative on Bitfinex (-0.21 bp) and large and positive on Bitstamp (2.83 bp). The difference in realized spreads between GDAX and Bitfinex is partly offset by the difference in maker fees. Liquidity suppliers on Bitfinex pay up to 0.1 bp (depending on their total trading volume during the previous 30 days) while GDAX charges no maker fee. The maker fees on Bitstamp, although higher than those of the other exchanges, are dwarfed by the large realized spreads. We thus conclude that supplying liquidity is more profitable on Bitstamp than on the other two exchanges.

Both the number and the volume of buyer-initiated and seller initiated trades are roughly equal on Bitfinex, as is shown by order imbalance measures close to zero. On Bitstamp and GDAX, on the other hand, the number and volume of buyer-initiated trades exceed the corresponding values for seller-initiated trades. This finding is consistent with our earlier result that the fraction of buyer-initiated trades is above 50% on Bitstamp and GDAX while it is slightly below 50% on Bitfinex.

The depth at the best quotes measures the average volume (in USD) that can be traded at the current best bid and ask prices. It is highest on GDAX (at 92,170 USD), followed by Bitfinex (43,590 UDS) and Bitstamp (16,690 USD). The lower right panel of figure 7 shows that this ranking is stable throughout the sample period. Thus, GDAX consistently and simultaneously has the lowest spreads and the highest depth while Bitstamp has the highest spreads and the lowest depth. The fee differences are much too low to outweigh these differences.

The upper panel of Table 7 reports correlations across liquidity measures, separately for each of the three exchanges. Quoted spreads are significantly positively correlated with effective spreads and negatively correlated with the depth at the best quotes on all three exchanges. The correlation with trading volume is positive on Bitfinex and GDAX but negative on Bitstamp.

The lower panel of Table 7 shows, for each liquidity measure, the correlation between the three exchanges. As before, the results are based on the synchronized 1-minute intervals. The correlations for dollar volume and the number

Table 6: Descriptive statistics for liquidity measures derived from the synced data set. QS, ES and PI refer to quoted spread, effective spread and price impact, respectively, all measured in basis points ([bp]). DV denotes the dollar volume and is reported in thousands of USD. OI and OIV refer to order imbalance and order imbalance volume, AvgD is the average BBO depth and is reported in thousand USD. Finally, numTA denotes the average number of transaction in a one-minute interval. The lower part of the table reports the mean and the median of the time delay between the first order book snapshot of an interval and the subsequent transaction (used in the calculation of the effective spread), and the time delay between this transaction and the next order book snapshot (used in the calculation of the price impact). Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

	Bitfinex		Bitstamp		GDAX	
	mean	median	mean	median	mean	median
QS [bp]	1.24	0.49	9.64	8.06	0.22	0.01
ES [bp]	1.09	0.15	7.94	7.05	0.51	0.01
PI [bp]	0.76	0	1.14	0.04	0.90	0
DV [1,000 USD]	303.71	145.04	102.30	38.66	119.80	49.35
OI	-0.00	-0.01	0.20	0.23	0.15	0.20
OIV	-0.01	-0.02	0.08	0.11	0.11	0.16
AvgD [1,000 USD]	43.59	22.87	16.69	7.23	92.17	71.37
numTA	82.41	57	32.61	22	56.93	33
Interim time between order book 1 (OB1), transaction (TA), order book 2 (OB2)						
OB1 - TA	4.19 sec	2 sec	7.25 sec	3 sec	4.71 sec	2.88 sec
TA - OB2	6.52 sec	6 sec	6.94 sec	5 sec	5.69 sec	5.44 sec

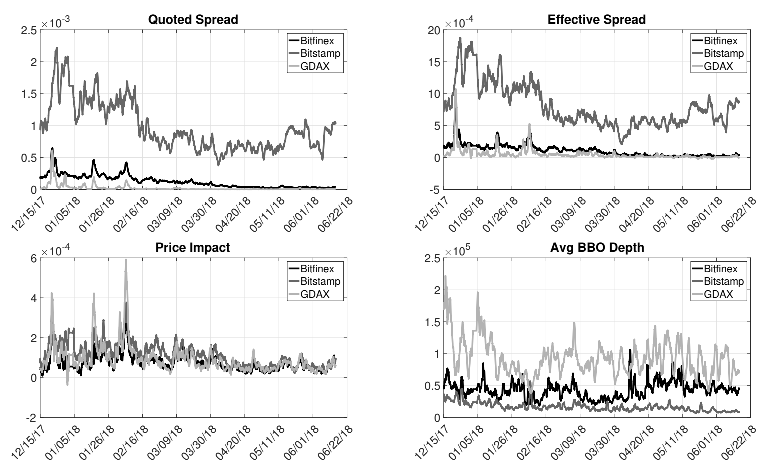

Figure 7: Liquidity measures quoted spread, effective spread, price impact and average BBO depth over time for all three exchanges. All measures are derived from synced one-minute-intervals, all subplots show a moving average of 1,440 one-minute intervals (24 hours). Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

Table 7: Correlation Analysis. The upper panel shows correlations of liquidity measures, based on one-minute intervals, for each of the three exchanges. QS, ES, PI, DV, OI, OIV, AvgD, and numTA refer to Quoted Spread, Effective Spread, Price Impact, Dollar Volume, Order Imbalance, volume weighted Order Imbalance, average BBO Depth, and the number of transaction per one-minute interval, respectively. "*" denotes statistical significance at the 5% level. The lower panel reports, for each liquidity measure, the pairwise correlations between the three exchanges; for instance, the value 0.13 in line 'BF/BS' in column 'QS' refers to the correlation of the quoted spread between Bitfinex and Bitstamp. Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

Bitfinex	QS	ES	PI	DV	OI	OIV	AvgD	numTA
ES	0.26*							
PI	0.06*	0.43*						
DV	0.20*	0.11*	0.07*					
OI	-0.02*	-0.01*	-0.01*	-0.04*				
OIV	-0.01*	-0.01*	-0.01*	-0.02*	0.80*			
AvgD	-0.05*	-0.02*	-0.00	0.12*	0.00	0.00		
numTA	0.28*	0.14*	0.08*	0.81*	-0.06*	-0.03*	0.03*	
Bitstamp								
ES	0.44*							
PI	0.01*	0.41*						
DV	-0.02*	-0.01*	0.08*					
OI	0.01*	0.02*	-0.01*	-0.12*				
OIV	0.01*	0.02*	-0.01*	-0.07*	0.69*			
AvgD	-0.05*	-0.03*	-0.00	0.22*	-0.03*	-0.01*		
numTA	-0.01*	-0.00	0.08*	0.73*	-0.12*	-0.06*	0.14*	
GDAX								
ES	0.09*							
PI	0.02*	0.74*						
DV	0.16*	0.10*	0.13*					
OI	-0.01*	0.00	-0.00	-0.09*				
OIV	-0.01*	0.00	-0.00	-0.07*	0.81*			
AvgD	-0.09*	-0.01*	-0.02*	0.09*	-0.01*	0.00		
numTA	0.17*	0.10*	0.15*	0.82*	-0.08*	-0.06*	0.00	
Cross exchange correlations								
BF/BS	0.13*	0.03*	0.03*	0.48*	0.22*	0.21*	0.01	0.70*
BF/GX	0.15*	0.03*	0.03*	0.57*	0.32*	0.28*	0.03*	0.70*
BS/GX	0.06*	0.00	0.03*	0.50*	0.25*	0.20*	0.03*	0.65*

of transactions are large and positive. Most other correlations are surprisingly low. These results strengthen the conclusion that the three trading venues are not integrated in terms of liquidity provision or activity.

The quoted bid-ask spread and quoted depth reported above only capture the liquidity that is available at the best quotes. For traders wishing to trade a quantity exceeding the depth at the best bid or ask, the liquidity further up or down the order book is important. We assess it by calculating, separately for each exchange, weighted average prices for buy and sell orders of specific sizes as described in section 3. We then calculate WAP, defined as the percentage difference between the weighted average prices for a buy and a sell order of equal size. The results are shown in Figure 8.

GDAX is the most liquid exchange irrespective of the order size considered.

This finding is consistent with the result, reported above, that GDAX offers the lowest quoted spreads and the highest depth. Bitstamp, on the other hand, is, by a large margin, the least liquid exchange. It is also apparent from the figures that there is a high correlation between the WAP measures for different trade sizes. Thus, considering the liquidity beyond the best quotes does not change our conclusions.

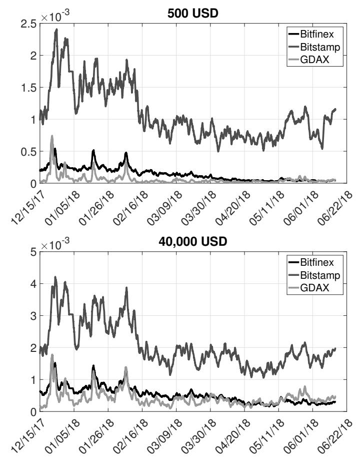

Figure 8: Percentage difference of Weighted Average Prices when buying / selling target volumes of 500, 2,000, 40,000, and 100,000 USD. The black / gray / light gray lines refer to data from Bitfinex / Bitstamp / GDAX respectively. Data is collected from Bitfinex, Bitstamp, and GDAX between 12/15/2017 and 06/15/2018.

6 Conclusion

The production of efficient prices is an important role of financial markets. We document that the decentralized and unregulated nature of cryptocurrency exchanges has led to transaction prices and limit orders that diverge from each other and perhaps the fundamental value. While we only study a short and recent period of time we also show that cryptocurrency prices (transaction and limit order) appear to become more integrated in the most recent period.

More efficient prices should lead to trust in cryptocurrency exchanges, increased trading (and higher welfare associated with gains from trade), and perhaps higher prices in the long-run. It will be interesting to document whether or not the laws of supply and demand and arbitrage constraints will be enough to align prices across multiple cryptocurrencies and exchanges. In order to remain viable in the long term the cryptocurrency market will have to develop an infrastructure that is similar to modern financial markets. A good first step may be the introduction of a central BTC depository or brokerage that would allow participants to trade on multiple exchanges without transferring BTC of fiat directly to each exchange. This may require exchanges to change their business models.

That cryptocurrency markets will remain mostly unregulated is unlikely. Policy makers around the world have implemented and are in the process of designing regulations that will impact cryptocurrencies. For instance, some jurisdictions are considering treating Bitcoin as a financial security. Other jurisdictions are limiting citizens to specific exchanges. Capital controls are also being implemented to limit purchases in local currencies. Law enforcement is tracking the illicit uses of Bitcoin. The future will certainly look different than the current regulatory environment and will impact the adoption of cryptocurrencies and exchanges.

Satoshi Nakamoto originally designed the Blockchain and Bitcoin as a response to fractional reserve banking and the financial crisis. Nakamoto embedded the following in the first block of the Blockchain: "The Times 03/Jan/2009 Chancellor on brink of second bailout for banks." Bitcoin was designed as an anonymous, decentralized system in competition with central bank controlled money supply. Exchanges partially circumvent this philosophy by introducing a non-anonymous centralized reserve of Bitcoin on each exchange. The next step towards a traditional banking system is ubiquitous Bitcoin lending. Crypto-exchanges require customers to fully reveal their identity. Customers of cryptocurrency exchanges must also transfer funds (fiat and cryptocurrency) and control over these funds to the exchange, similar to the banking system. Usurping the financial system may be harder and less efficient than originally thought.

Future work should be focused on studying cross-asset (for example, Ethereum or Ripple) prices and price discovery and extending the sample period. New studies could also focus on infrastructure and exchange rule changes and regulatory events. Finally, establishing a tighter link between on-chain and off-chain activity is an important avenue of future inquiry.

References

Abadi, Joseph and Markus Brunnermeier (2018). *Blockchain Economics*. Tech. rep. mimeo Princeton University. URL: `https://scholar.princeton.edu/sites/default/files/markus/files/blockchain_paper_v3g.pdf`.

Acharya, Viral V. and Lasse Heje Pedersen (2005). "Asset pricing with liquidity risk". In: *Journal of Financial Economics* 77.2, pp. 375–410.

Amihud, Yakov and Haim Mendelson (1986). "Asset pricing and the bid-ask spread". In: *Journal of Financial Economics* 17.2, pp. 223–249.

Amihud, Yakov, Haim Mendelson, and Beni Lauterbach (1997). "Market microstructure and securities values: Evidence from the Tel Aviv Stock Exchange". In: *Journal of Financial Economics* 45.3, pp. 365–390.

Aoyagi, AJun and Daisuke Adachi (2018). "Fundamental Values of Cryptocurrencies and Blockchain Technology". In: *Available at SSRN*. URL: `https://ssrn.com/abstract=3132235`.

Bariviera, Aurelio F. (2017). "The inefficiency of Bitcoin revisited: A dynamic approach". In: *Economics Letters* 161, pp. 1–4.

Bariviera, Aurelio F. et al. (2017). "Some stylized facts of the Bitcoin market". In: *Physica A: Statistical Mechanics and its Applications* 484, pp. 82–90.

Baur, Dirk G., Thomas Dimpfl, and Konstantin Kuck (2018). "Bitcoin, gold and the US dollar – A replication and extension". In: *Finance Research Letters* 25, pp. 103–110.

Beber, Alessandro and Marco Pagano (2013). "Short-Selling Bans Around the World: Evidence from the 2007–09 Crisis". In: *The Journal of Finance* 68.1, pp. 343–381.

Bessembinder, Hendrik and Kumar Venkataraman (2010). "Bid–Ask Spreads". In: *Encyclopedia of Quantitative Finance*. American Cancer Society.

Boehmer, Ekkehart, Kingsley Fong, and Julie Wu (2015). "International Evidence on Algorithmic Trading". In: *https://ssrn.com/abstract=2022034* 2015.

Boehmer, Ekkehart, Charles Jones, and Xiaoyan Zhang (2013). "Shackling Short Sellers: The 2008 Shorting Ban". In: *The Review of Financial Studies* 26.6, pp. 1363–1400.

Bouri, Elie et al. (2017). "On the hedge and safe haven properties of Bitcoin: Is it really more than a diversifier?" In: *Finance Research Letters* 20, pp. 192–198.

Brauneis, Alexander and Roland Mestel (2018a). "Cryptocurrency-portfolios in a mean-variance framework". In: *Finance Research Letters, forthcoming*.

— (2018b). "Price Discovery of Cryptocurrencies: Bitcoin and beyond". In: *Economic Letters* 165, pp. 58–61.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan (2014). "High-Frequency Trading and Price Discovery". In: *The Review of Financial Studies* 27.8, pp. 2267–2306.

— (2017). "High frequency trading and the 2008 short-sale ban". In: *Journal of Financial Economics* 124.1, pp. 22–42.

Brunnermeier, Markus and Lasse Heje Pedersen (2008). "Market liquidity and funding liquidity". In: *The Review of Financial Studies* 22.6, pp. 2201–2238.

Chaboud, Alain P. et al. (2014). "Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market". In: *The Journal of Finance* 69.5, pp. 2045–2084.

Chakrabarty, Bidisha, Jianning Huang, and Pankaj K. Jain (2018). "Exchange Competition with Levelled Speed". In: *Available at SSRN.* URL: https://ssrn.com/abstract=3197088.

Colliard, Jean-Edouard and Thierry Foucault (2012). "Trading Fees and Efficiency in Limit Order Markets". In: *Available at SSRN.* URL: https://ssrn.com/abstract=1831146.

Cong, Lin William and Zhiguo He (2018). *Blockchain disruption and smart contracts.* Tech. rep. National Bureau of Economic Research.

Corbet, Shaen et al. (2018). "Exploring the dynamic relationships between cryptocurrencies and other financial assets". In: *Economics Letters* 165, pp. 28–34.

Degryse, Hans (2009). "Competition between financial markets in Europe: what can be expected from MiFID?" In: *Financial Markets and Portfolio Management* 23.1, pp. 93–103.

Degryse, Hans, Frank de Jong, and Vincent van Kervel (2015). "The Impact of Dark Trading and Visible Fragmentation on Market Quality". In: *Review of Finance* 19.4, pp. 1587–1622.

Dimpfl, Thomas (2017). "Bitcoin Market Microstructure". In: *working paper*, pp. 1–24.

Dyhrberg, Anne Haubo (2016a). "Bitcoin, gold and the dollar – A GARCH volatility analysis". In: *Finance Research Letters* 16, pp. 85–92.

— (2016b). "Hedging capabilities of bitcoin. Is it the virtual gold?" In: *Finance Research Letters* 16, pp. 139–144.

ESMA (2014). *High-frequency trading activity in EU equity markets.* Paris. URL: https://www.esma.europa.eu/sites/default/files/library/2015/11/esma20141_-_hft_activity_in_eu_equity_markets.pdf.

Fink, Christopher and Thomas Johann (2014). "Bitcoin Markets". In: *Available at SSRN.* URL: https://ssrn.com/abstract=2408396.

Foley, Sean, Jonathan Karlsen, and Tālis J Putniņš (2018). "Sex, Drugs, and Bitcoin: How Much Illegal Activity is Financed Through Cryptocurrencies?" In: *Available at SSRN.* URL: https://ssrn.com/abstract=3102645.

Foucault, Thierry et al. (2013). *Market liquidity: theory, evidence, and policy.* Oxford University Press.

Gandal, Neil et al. (2018). "Price manipulation in the Bitcoin ecosystem". In: *Journal of Monetary Economics*, pp. 86–96.

Glosten, Lawrence R. (1987). "Components of the Bid-Ask Spread and the Statistical Properties of Transaction Prices". In: *The Journal of Finance* 43.5, pp. 1293–1307.

Glosten, Lawrence R. and Paul R. Milgrom (1985). "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders". In: *Journal of Financial Economics* 14.1, pp. 71–100.

Gresse, Carole (2017). "Effects of lit and dark market fragmentation on liquidity". In: *Journal of Financial Markets* 35, pp. 1–20.

Hansch, Oliver, Narayan Y Naik, and S Viswanathan (1999). "Preferencing, internalization, best execution, and dealer profits". In: *The Journal of Finance* 54.5, pp. 1799–1828.

Harris, Lawrence (1999). "Trading in pennies: a survey of the issues". In: *Working Paper*. URL: https://ssrn.com/abstract=1831146.

Hasbrouck, Joel and Gideon Saar (2013). "Low-latency trading". In: *Journal of Financial Markets* 16.4. High-Frequency Trading, pp. 646–679.

Haslag, Peter and Matthew Ringgenberg (2016). "The Causal Impact of Market Fragmentation on Liquidity". In: *Available at SSRN*. URL: https://ssrn.com/abstract=2591715.

Hendershott, Terrence, Charles Jones, and Albert Menkveld (2011). "Does Algorithmic Trading Improve Liquidity?" In: *The Journal of Finance* 66.1, pp. 1–33.

Hengelbrock, Jördis and Erik Theissen (2009). "Fourteen at One Blow: The Market Entry of Turquoise". In: *Working Paper*.

Holden, C.W., S. Jacobsen, and A. Subrahmanyam (2014). *The Empirical Analysis of Liquidity*. Foundations and Trends in Finance. now Publishers Incorporated.

Hu, Albert, Christine A Parlour, and Uday Rajan (2018). "Cryptocurrencies: Stylized Facts on a New Investible Instrument". In: *Available at SSRN*. URL: https://ssrn.com/abstract=3182113.

Huang, Jennifer and Jiang Wang (2010). "Market liquidity, asset prices, and welfare". In: *Journal of Financial Economics* 95.1, pp. 107–127.

Huang, Roger and Hans Stoll (1996). "Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE". In: *Journal of Financial Economics* 41.3, pp. 313–357.

Kroeger, Alexander and Asani Sarkar (2017). "Law of one bitcoin price?" In: *Federal Reserve Bank of Philadelphia*. URL: www.iaes.org/.../Slides_Sarkar_Bitcoin_Philadelphia_ARK-1.pdf.

Macey, Jonathan R and Maureen O'Hara (1997). "The law and economics of best execution". In: *Journal of Financial Intermediation* 6.3, pp. 188–223.

Makarov, Igor and Antoinette Schoar (2018). "Trading and Arbitrage in Cryptocurrency Markets". In: *Available at SSRN*. URL: https://ssrn.com/abstract=3171204.

Malinova, Katya and Andreas Park (2015). "Subsidizing Liquidity: The Impact of Make/Take Fees on Market Quality". In: *The Journal of Finance* 70.2, pp. 509–536.

Malinova, Katya, Andreas Park, and Ryan Riordan (2017). "Do retail traders suffer from high frequency traders?" In: *Available at SSRN*. URL: https://ssrn.com/abstract=2183806.

Nadarajah, Saralees and Jeffrey Chu (2017). "On the inefficiency of Bitcoin". In: *Economics Letters* 150, pp. 6–9.

Nakamoto, Satoshi (2008). "Bitcoin: A Peer-to-Peer Electronic Cash System". In: URL: https://bitcoin.org/bitcoin.pdf.

O'Hara, Maureen and Mao Ye (2011). "Is market fragmentation harming market quality?" In: *Journal of Financial Economics* 100.3, pp. 459–474.

Pagnotta, Emiliano and Andrea Buraschi (2018). “An Equilibrium Valuation of Bitcoin and Decentralized Network Assets”. In: *Available at SSRN*. URL: https://ssrn.com/abstract=3142022.

Pastor, Lubos and Robert F Stambaugh (2003). “Liquidity Risk and Expected Stock Returns”. In: *Journal of Political Economy* 111.3, pp. 642–685.

Rindi, Barbara and Ingrid M Werner (2017). “US Tick Size Pilot”. In: *Available at SSRN*. URL: https://ssrn.com/abstract=3041644.

Riordan, Ryan and Andreas Storkenmaier (2012). “Latency, liquidity and price discovery”. In: *Journal of Financial Markets* 15.4, pp. 416–437.

Schilling, Linda and Harald Uhlig (2018). “Some Simple Bitcoin Economics”. In: *Available at SSRN*. URL: https://ssrn.com/abstract=3155310.

Shi, Shimeng (2018). “The Impact of Futures Trading on Intraday Spot Volatility and Liquidity: Evidence from Bitcoin Market”. In: *working paper*, pp. 1–14.

Shleifer, Andrei and Robert W Vishny (1997). “The limits of arbitrage”. In: *The Journal of Finance* 52.1, pp. 35–55.

Urquhart, Andrew (2016). “The inefficiency of Bitcoin”. In: *Economics Letters* 148, pp. 80–82.

Urquhart, Andrew and Hanxiong Zhang (2018). “Is Bitcoin a Hedge or Safe-Haven for Currencies? An Intraday Analysis”. In: *Available at SSRN*. URL: https://ssrn.com/abstract=3114108.